

LOKESH DHAKAR

+91-9755786633 | lokeshdhakar1404@gmail.com | LinkedIn: linkedin.com/in/iamlokesdhakar | GitHub: github.com/iamlokesdhakar | Portfolio

Professional Summary

Founding Engineer and product-focused Full-Stack Engineer with 2.5+ years of experience building production systems in early-stage startups. First engineer hired at Axy, helping rebuild a scientific knowledge graph platform from the ground up across frontend, backend, distributed systems, and applied AI. Hands-on ownership of scalable systems supporting millions of graph entities, from product requirements and architecture to implementation and production.

Work Experience

Founding Engineer | Axy Inc

Aug 2025 – Present

Remote

- Joined as the first engineer hired and helped rebuild the product from an MVP into a production scientific knowledge graph platform; co-designed the Turborepo monorepo, BFF/API gateway, service boundaries, shared packages, and backend architecture while working closely with the CTO, designer, and scientists on product requirements.
- Built and scaled data systems supporting millions of scientific knowledge graph entities across PostgreSQL, MongoDB, ArangoDB, Qdrant, and Redis, powering graph exploration, research workflows, document management, and semantic search.
- Designed and implemented the end-to-end submission workflow for researchers to draft graph changes, review and approve submissions, detect conflicts, maintain mutation history, and safely merge changes into the main graph using a delta-based model.
- Moved cross-service synchronization and long-running submission merges out of synchronous API requests using AWS Lambda, SNS, and SQS; added dead-letter handling and reconciliation checks to keep graph, search, vector, and document services consistent.
- Built a multi-stage retrieval pipeline using dense and sparse search in Qdrant, Reciprocal Rank Fusion, and reranking, improving retrieval quality from roughly 0.5 baseline to 0.7–0.8 across scientific BEIR evaluations using Recall@K, NDCG@10, and MRR.
- Owned major parts of the applied-AI pipeline including embeddings, hybrid retrieval, query rewriting, reranking, context selection, agent/tool workflows, chunking, and evaluation; also built a Python prototype for converting scientific papers into structured knowledge graph entities.

Full Stack Developer | AlphaBI

Feb 2024 – Jul 2025

Surat, Gujarat

- Rebuilt major parts of an IoT monitoring dashboard by developing 10+ production pages from Figma designs and helped migrate backend APIs from Express to Fastify with schema validation and serialization, achieving approximately 1.5× API throughput.
- Independently built and shipped an internal HRMS covering authentication, attendance, punch-in/out, and employee workflows within a two-week sprint using AI-assisted development workflows.
- Developed and integrated an internal NPM logging package across three products using Axios interceptors to capture contextual API logs alongside user-reported issues, improving production debugging and root-cause analysis.
- Added unit, integration, and end-to-end tests using Jest, Playwright, and Cypress, reaching approximately 75% coverage on critical flows while improving structured logging and error observability across production services.

Technical Skills

Languages & Frameworks: TypeScript, JavaScript, Python, SQL, React, Next.js, Node.js, Fastify, Express.js

Data & Search: PostgreSQL, MongoDB, ArcadeDB, Qdrant Vector Database, Redis

AI & Cloud: RAG, LLM Integration, Embeddings, Hybrid Search, Reranking, Retrieval Evaluation, AWS Lambda, SQS, SNS, S3

Tools & Platform: Tailwind CSS, Zustand, TanStack Query, Zod, Docker, Turborepo, GitHub Actions, Jest, Vitest, Playwright, WorkOS, OpenFGA

Education

B.Tech in Computer Science | RGPV University

2020 – 2024

Sagar Institute of Research Technology and Science (SIRTS), Bhopal

- CGPA: 8.4/10 — Coursework: Data Structures and Algorithms, DBMS, OS, CN